

IN THE SPECIFICATION:

Please replace the second full paragraph of specification page 6 with the following replacement paragraph:

Generally, a PCPI is an *image* (typically read-only) of a file system at a point in time, which is stored on the same primary storage device as is the active file system and is accessible by users of the active file system. Note, that by “active file system” it is meant the file system to which current input/output (I/O) operations are being directed. Each time a PCPI occurs, the old active file system becomes the new PCPI, and the new active file system carries on, recording any new changes. A set number of PCPIs may be retained depending upon various time-based and other criteria. The PCPI process is described in further detail in United States Patent Application Serial No. 09/932,578, entitled INSTANT SNAPSHOT by Blake Lewis *et al.*, now issued as U.S. Patent No. 7,454,445 on November 18, 2008, which is hereby incorporated by reference as though fully set forth herein

Please replace the first full paragraph of specification page 8 with the following replacement paragraph:

Approaches to volume-based remote mirroring of PCPIs are described in detail in commonly owned U.S. Patent Application Serial No. 09/127,497, entitled FILE SYSTEM IMAGE TRANSFER by Steven Kleiman, *et al.*, now issued as U.S. Patent No. 6,604,118 on August 5, 2003, and U.S. Patent Application Serial No. 09/426,409, entitled FILE SYSTEM IMAGE TRANSFER BETWEEN DISSIMILAR FILE SYSTEMS by Steven Kleiman, *et al.*, now issued as U.S. Patent No. 6,574,591 on June 3, 2003, both of which are expressly incorporated herein by reference.

Please replace the first full paragraph of specification page 9 with the following replacement paragraph:

One such sub-organization of a volume is the well-known qtree. Qtrees, as implemented on an exemplary storage system such as described herein, are subtrees in a volume's file system. One key feature of qtrees is that, given a particular qtree, any file or directory in the system can be quickly tested for membership in that qtree, so they serve as a good way to organize the file system into discrete data sets. The use of qtrees as a source and destination for replicated data may be desirable. An approach to remote asynchronous mirroring of a qtree is described in U.S. Patent Application Serial No. 10/100,967 entitled SYSTEM AND METHOD FOR DETERMINING CHANGES IN TWO SNAPSHOTS AND FOR TRANSMITTING CHANGES TO A DESTINATION SNAPSHOT, by Michael L. Federwisch, et al., now issued as U.S Patent No. 6,993,539 on January 31, 2006, the teachings of which are incorporated herein by reference.

Please replace the paragraph bridging specification pages 17 and 18 with the following replacement paragraph:

In an illustrative embodiment, the disks 660 are arranged into a plurality of volumes (for example, source volumes 614 and destination volumes 616), in which each volume has a file system associated therewith. The volumes each include one or more disks 660. In one embodiment, the physical disks 660 are configured into RAID groups so that some disks store striped data and some disks store separate parity for the data, in accordance with a preferred RAID 4 configuration. However, other configurations (e.g. RAID 5 having distributed parity across stripes) are also contemplated. In this embodiment, a minimum of one parity disk and one data disk is employed. However, a typical implementation may include three data and one parity disk per RAID group, and a multiplicity of RAID groups per volume. Information in the volumes may be organized as files, directories and virtual disks (vdisks). That is, to facilitate access to the disks, the

destination storage operating system 700 implements a write-anywhere file system that logically organizes the information as a hierarchical structure of directory, file and vdisk objects (hereinafter “directories”, “files” and “vdisks”) on the disks. A vdisk is a special file type that is translated into an emulated disk or logical unit number (lun) as viewed by a storage area network (SAN) client, which may implement a block-based protocol for storage organization. Each “on-disk” file may be implemented as set of disk blocks configured to store information, such as data, whereas the directory may be implemented as a specially formatted file in which names and links to other files and directories are stored. Vdisks are further described in U.S. Patent Application Serial No. 10/216,453, entitled **STORAGE VIRTUALIZATION BY LAYERING VIRTUAL DISK OBJECTS ON A FILE SYSTEM**, by Vijayan Rajan, et al., now issued as U.S. Patent No. 7,107,385 on September 12, 2006, the teachings of which are hereby incorporated by reference. As will be made more clear with reference to the storage operating system 700 below, the exemplary destination filer 700 may be characterized more particularly as a multi-protocol storage appliance, capable of receiving storage requests according to a plurality of protocols and performing storage service operations according to either a file-based or block-based organization.

Please replace the first full paragraph on specification page 18 with the following replacement paragraph:

In the illustrative embodiment, the source volume 614 is mirrored to the destination volume using a qtree-based mirroring technique, such as that described in the above-incorporated patent application entitled **SYSTEM AND METHOD FOR DETERMINING CHANGES IN TWO SNAPSHOTS AND FOR TRANSMITTING CHANGES TO A DESTINATION SNAPSHOT**, by Michael L. Federwisch, et al. However, it should be noted that the technique of the present invention may be utilized with any acceptable mirroring technique that may be suitably modified to utilize the teachings of the present invention. To facilitate the transfer of replica data from the source 610 to the destination

612 via the network 618, a formatted data stream 690 is employed. While it is contemplated that a variety of data transmission schemes and formats can be employed according to various alternate embodiments of this invention, and exemplary data format and related concepts are described in commonly owned U.S. patent application Serial No. 10/777,978, entitled **SYSTEM-INDEPENDENT DATA FORMAT IN A MIR-**

RORED STORAGE SYSTEM ENVIRONMENT AND METHOD FOR USING THE SAME, by Stephen L. Manley, now issued as U.S. Patent No. 7,478,101 on January 13, 2009, the teachings of which are expressly incorporated herein by reference.

Please replace the paragraph bridging specification pages 18 and 19 with the following replacement paragraph:

The destination storage system 612, as the repository of backup data, organized a set of PCPIs, interfaces with an administrator's management client 692. This can be any acceptable computer or terminal console that runs or manipulates a software backup management application 693 and that (in this example) provides the user/administrator with a graphical user interface (GUI) 694 for displaying and manipulating data. The graphical user interface in this embodiment employs a web-capable browser application 695 that enables data to be displayed and controlled via conventional web pages 696. The web pages are generated using conventional Hypertext Markup Language (HTML) script 697, an engine for which is provided in the backup management application 693. The administrator's management client, the user interface, and its communication with the file server are described, by way of further background, in U.S. Patent Application Serial No. 09/862,949, entitled **SYSTEM AND METHOD FOR CONSOLIDATED REPORTING OF CHARACTERISTICS FOR A GROUP OF FILE SYSTEMS**, by Brian M. Hackworth, published as U.S. Patent Publication No. 2002/0175938 on November 28, 2002, the teachings of which are expressly incorporated herein by reference. In an exemplary commercial implementation, the system software (for use with compatible third-

party hardware) is available from Network Appliance, Inc. of Sunnyvale, CA as the Data Fabric Manager™.

Please replace the paragraph bridging specification pages 23 and 24 with the following replacement paragraph:

By way of further background, the general principles of remote asynchronous mirroring of a volume or sub-volume/qtree data set from a source to a destination are described in detail in the above-incorporated U.S. Patent application Serial No. 10/100,950, entitled SYSTEM AND METHOD FOR ASYNCHRONOUS MIRRORING OF SNAPSHOTS AT A DESTINATION USING A PURGATORY DIRECTORY AND INODE MAPPING by Stephen L. Manley, *et al.*, now issued as U.S. Patent No. 7,225,204 on May 29, 2007, and other related applications incorporated hereinabove. The illustrative embodiment of the invention employs the techniques described in these above-incorporated patent applications to receive and store, as a replica data set of the source, on the destination side, a data stream generated by the source. The source data may be organized in a variety of manners and appropriate mechanisms on the source side (in its replication agent 613) are used to generate the formatted data stream 690 in accordance with the illustrative embodiment of this invention. In one embodiment, both the source system and the destination system include respective pipelines (described further below) for generating and decoding the formatted data stream.